Subshifts and Logic: Back and Forth

Emmanuel Jeandel[a], Guillaume Theyssier[b]

[a] *Laboratoire d'informatique fondamentale de Marseille (LIF)*
Aix-Marseille Université, CNRS
39 rue Joliot-Curie, 13453 Marseille Cedex 13, FRANCE
[b] *LAMA (Université de Savoie, CNRS)*
Campus Scientifique, 73376 Le Bourget-du-lac cedex FRANCE

Abstract

We study the Monadic Second Order (MSO) Hierarchy over colourings of the discrete plane, and draw links between classes of formula and classes of subshifts. We give a characterization of existential MSO in terms of projections of tilings, and of universal sentences in terms of combinations of "pattern counting" subshifts. Conversely, we characterise logic fragments corresponding to various classes of subshifts (subshifts of finite type, sofic subshifts, all subshifts). Finally, we show by a separation result how the situation here is different from the case of tiling pictures studied earlier by Giammarresi et al.

Key words: Symbolic Dynamics, Model Theory, Tilings

1. Introduction

There is a close connection between words and monadic second-order (MSO) logic. Büchi and Elgot proved for finite words that MSO-formulas correspond exactly to regular languages. This relationship was developed for other classes of labeled graphs; trees or infinite words enjoy a similar connection. See [1, 2] for a survey of existing results. Colorings of the entire plane, i.e tilings, represent a natural generalization of biinfinite words to higher dimensions, and as such enjoy similar properties. We plan to study in this paper tilings for the point of view of monadic second-order logic.

Tilings and logic have a shared history. The introduction of tilings can be traced back to Hao Wang [3], who introduced his celebrated tiles to study the (un)decidability of the $\forall\exists\forall$ fragment of first order logic. The undecidability of the domino problem by his PhD Student Berger [4] lead then to the undecidability of this fragment [5]. Seese [6, 7] used the domino problem to prove that graphs with a decidable MSO theory have a bounded tree width. Makowsky[8, 9] used the construction by Robinson [10] to give the first example of a finitely axiomatizable super-stable theory that is super-stable. More recently, Oger [11] gave generalizations of classical results on tilings to locally finite relational structures. See the survey [12] for more details.

Previously, a finite variant of tilings, called tiling pictures, was studied [13, 14]. Tiling pictures correspond to colorings of a *finite* region of the plane, this region being bordered by special '`#`' symbols. It is proven for this particular model that language recognized by EMSO-formulas correspond exactly to so-called finite tiling systems, i.e. projections of finite tilings.

The equivalent of finite tiling systems for infinite pictures are so-called *sofic subshifts* [15]. A *sofic subshift* represents intuitively local properties and ensures that every point of the plane behaves in the same way. As a consequence, there is no general way to enforce that some specific color, say , appears at least once. Hence, some simple first-order existential formulas have no equivalent as sofic subshift (and even subshift). This is where the border of `#` for finite pictures play an important role: Without such a border, results on finite pictures would also stumble on this issue. See [16] for similar results on finite pictures without borders.

Email addresses: `Emmanuel.Jeandel@lif.univ-mrs.fr` (Emmanuel Jeandel), `guillaume.theyssier@univ-savoie.fr` (Guillaume Theyssier)

We deal primarily in this article with subshifts. See [17] for other acceptance conditions (what we called subshifts of finite type correspond to A-acceptance in this paper).

Finally, note that all decision problems in our context are non-trivial : To decide if a universal first-order formula is satisfiable (the domino problem, presented earlier) is not recursive. Worse, it is Σ_1^1-hard to decide if a tiling of the plane exists where some given color appears infinitely often [18, 17]. As a consequence, the satisfiability of MSO-formulas is at least Σ_1^1-hard.

In this paper, we will prove how various classes of formula correspond to well known classes of subshifts. Some of the results of this paper were already presented in [19].

2. Symbolic Spaces and Logic

2.1. Configurations

Consider the discrete lattice $\mathbb{Z}^2$. For any finite set Q, a Q-configuration is a function from $\mathbb{Z}^2$ to Q. Q may be seen as a set of *colors* or *states*. An element of $\mathbb{Z}^2$ will be called a *cell*. A configuration will usually be denoted C, M or N.

Fig. 1 shows an example of two different configurations of $\mathbb{Z}^2$ over a set Q of 5 colors. As a configuration is infinite, only a finite fragment of the configurations is represented in the figure. The reader has to use his imagination to decide what colors do appear in the rest of the configuration. We choose not to represent which cell of the picture is the origin $(0, 0)$. This will indeed be of no importance as we use only translation invariant properties.

For any $z \in \mathbb{Z}^2$ we denote by σ_z the *shift* map of vector z, *i.e.* the function from Q-configurations to Q-configurations such that for all $C \in Q^{\mathbb{Z}^2}$:

$$\forall z' \in \mathbb{Z}^2,\ \sigma_z(C)(z') = C(z' - z).$$

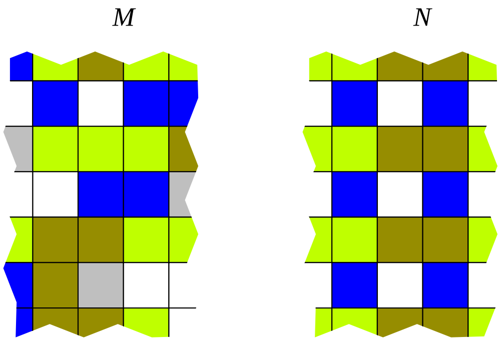

Figure 1: Two configurations

A *pattern* is a partial configuration. A pattern $P : X \to Q$ where $X \subseteq \mathbb{Z}^2$ occurs in $C \in Q^{\mathbb{Z}^2}$ at position z_0 if

$$\forall z \in X,\ C(z_0 + z) = P(z).$$

We say that P occurs in C if it occurs at some position in C. As an example the pattern P of Fig 2 occurs in the configuration M but not in N (or more accurately not on the finite fragment of N depicted in the figure). A finite pattern is a partial configuration of finite domain. All patterns in the following will be finite. The *language* $\mathcal{L}(C)$ of a configuration C is the set of finite patterns that occur in C. We naturally extend this notion to sets of configurations.

A *subshift* is a natural concept that captures both the notion of *uniformity* and *locality*: the only description "available" from a configuration C is the finite patterns it contains, that is $\mathcal{L}(C)$. Given a set $\mathcal{F}$ of patterns, let $X_{\mathcal{F}}$ be the set of all configurations where no patterns of $\mathcal{F}$ occurs.

$$X_{\mathcal{F}} = \{C | \mathcal{L}(C) \cap \mathcal{F} = \emptyset\}$$

$\mathcal{F}$ is usually called the set of forbidden patterns or the *forbidden language*. A set of the form $X_{\mathcal{F}}$ is called a *subshift*.

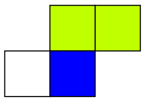

Figure 2: A pattern P. P appears in M but presumably not in N

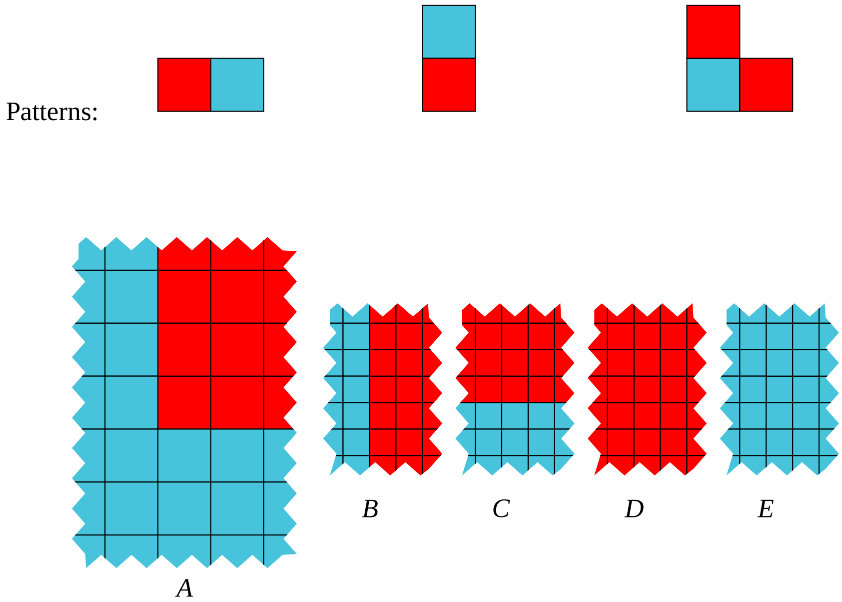

Figure 3: A (finite) set of forbidden patterns $\mathcal{F}$ and the tilings it generates

A subshift can be equivalentely defined by topology considerations. Endow the set of configurations $Q^{\mathbb{Z}^2}$ with the product topology: A sequence $(C_n)_{n\in\mathbb{N}}$ of configurations converges to a configuration C if the sequence ultimately agree with C on every $z \in \mathbb{Z}^2$. Then a subshift is a closed subset of $Q^{\mathbb{Z}^2}$ also closed by shift maps.

Example 1. Consider the three forbidden patterns of figure 3. The first one says that we cannot find a point at the left of a point. This can be interpreted as follows: every time we find a point, then all the points at the right of it are also . With the second forbidden pattern, we deduce that every time we find a point, then the entire quarter of plane on the above right of it is also filled with points. The third pattern ensures us that every configuration contains at most one quarter of plane of color : if it contains two such quarters of plane, then there must be a bigger quarter of plane that contains both.

Hence a typical configuration looks like A. Other possible configurations are B, C, D, E. They correspond to extremal situations where the corner of the quarter of plane is situated respectively at $(0, -\infty)$, $(-\infty, 0)$, $(-\infty, -\infty)$ et $(+\infty, +\infty)$

Example 2. Consider the set of colors $\{■, □\}$ and $\mathcal{F}$ to be the set of patterns that contains two ■ points or more.

Then $X_{\mathcal{F}}$ contains configurations with at most one ■ point. Up to shift, $X_{\mathcal{F}}$ contains then two configurations: the all □-one, and one where only one point is ■ and all others are □.

A *subshift of finite type* (or *tiling*) correspond to a finite set $\mathcal{F}$: it is the set of configurations C such that no pattern in $\mathcal{F}$ occurs in C. If all patterns of $\mathcal{F}$ are of diameter n, this means that we only have to see a configuration through a window of size n to know if it is a tiling, hence the locality. Example 1 is a subshift of finite type. It can be proven that Example 2 is not.

Given two state sets Q_1 and Q_2, a projection is a map $\pi : Q_1 \to Q_2$. We naturally extend it to $\pi : Q_1^{\mathbb{Z}^2} \to Q_2^{\mathbb{Z}^2}$ by $\pi(C)(z) = \pi(C(z))$. A *sofic subshift* of state set Q_2 is the image by some projection π of some subshift of finite type of state set Q_1. It is also a subshift (clearly closed by shift maps, and topologically closed because projections are continuous maps on a compact space). A sofic subshift is a natural object in tiling theory, although quite never mentioned explicitly. It represents the concept of *decoration*: some of the tiles we assemble to obtain the tilings may be decorated, but we forgot the decoration when we observe the tiling.

Example 3. Consider the following variant of Example 1: tilings are exactly the same except that the corner of the quarter of plane in A is of a different color $\color{lime}{\blacksquare}$. It is easy to see that this variant defines a subshift of finite type X (with a few more forbidden patterns).

Now consider the following map:

$$\pi : \begin{array}{ccc} \color{cyan}{\blacksquare} & \mapsto & \square \\ \color{red}{\blacksquare} & \mapsto & \square \\ \color{lime}{\blacksquare} & \mapsto & \blacksquare \end{array}$$

Then B, C, D, E will become under π of color $\square$, while A will become a configuration with exactly one $\blacksquare$, all other points being $\square$.

As a consequence, $\pi(X)$ is exactly Example 2. Example 2 is thus a sofic subshift.

2.2. *Structures*

A configuration will be seen in this article as an infinite structure. The signature τ contains four unary maps North, South, East, West and a predicate P_c for each color $c \in Q$.

A configuration M will be seen as a structure $\mathfrak{M}$ in the following way:

- The elements of $\mathfrak{M}$ are the points of $\mathbb{Z}^2$.
- North is interpreted by $\mathsf{North}^{\mathfrak{M}}((x,y)) = (x, y+1)$, East is interpreted by $\mathsf{East}^{\mathfrak{M}}((x,y)) = (x+1, y)$. $\mathsf{South}^{\mathfrak{M}}$ and $\mathsf{West}^{\mathfrak{M}}$ are interpreted similarly
- $P_c^{\mathfrak{M}}((x,y))$ is true if and only if the point at coordinate (x,y) is of color c, that is if $M(x,y) = c$.

As an example, the configuration M of Fig. 1 has three consecutive cells with the color $\color{yellowgreen}{\blacksquare}$. That is, the following formula is true:

$$\mathfrak{M} \models \exists z, P_{\color{yellowgreen}{\blacksquare}}(z) \land P_{\color{yellowgreen}{\blacksquare}}(\mathsf{East}(z)) \land P_{\color{yellowgreen}{\blacksquare}}(\mathsf{East}(\mathsf{East}(z)))$$

As another example, the following formula states that the configuration has a vertical period of 2 (the color in the cell (x,y) is the same as the color in the cell $(x, y+2)$). The formula is false in the structure $\mathfrak{M}$ and true in the structure $\mathfrak{N}$ (if the reader chose to color the cells of N not shown in the picture correctly):

$$\forall z, \begin{cases} P_{\color{olive}{\blacksquare}}(z) \implies P_{\color{olive}{\blacksquare}}(\mathsf{North}(\mathsf{North}(z))) \\ P_{\square}(z) \implies P_{\square}(\mathsf{North}(\mathsf{North}(z))) \\ P_{\color{blue}{\blacksquare}}(z) \implies P_{\color{blue}{\blacksquare}}(\mathsf{North}(\mathsf{North}(z))) \\ P_{\color{yellowgreen}{\blacksquare}}(z) \implies P_{\color{yellowgreen}{\blacksquare}}(\mathsf{North}(\mathsf{North}(z))) \\ P_{\color{gray}{\blacksquare}}(z) \implies P_{\color{gray}{\blacksquare}}(\mathsf{North}(\mathsf{North}(z))) \end{cases}$$

2.3. *Monadic Second-Order Logic*

This paper studies connection between subshifts (seen as structures as explained above) and monadic second order sentences. First order variables (x, y, z, ...) are interpreted as points of $\mathbb{Z}^2$ and (monadic) second order variables (X, Y, Z, ...) as subsets of $\mathbb{Z}^2$.

Monadic second order formulas are defined as follows:

- a term is either a first-order variable or a function (South, North, East, West) applied to a term ;
- atomic formulas are of the form $t_1 = t_2$ or $X(t_1)$ where t_1 and t_2 are terms and X is either a second order variable or a color predicate ;
- formulas are build up from atomic formulas by means of boolean connectives and quantifiers $\exists$ and $\forall$ (which can be applied either to first-order variables or second order variables).

A formula is *closed* if no variable occurs free in it. A formula is FO if no second-order quantifier occurs in it. A formula is EMSO if it is of the form

$$\exists X_1, \ldots, \exists X_n, \phi(X)$$

where ϕ is FO. Given a formula $\phi(X_1, \ldots, X_n)$ with no free first-order variable and having only $X_1, \ldots, X_n$ as free second-order variables, a configuration M together with subsets $E_1, \ldots, E_n$ is a model of $\phi(X_1, \ldots, X_n)$, denoted

$$(M, E_1, \ldots, E_n) \models \phi(X_1, \ldots, X_n),$$

if ϕ is satisfied (in the usual sense) when M is interpreted as $\mathfrak{M}$ (see previous section) and E_i interprets X_i.

2.4. *Definability*

This paper studies the following problems: Given a formula ϕ of some logic, what can be said of the configurations that satisfy ϕ? Conversely, given a subshift, what kind of formula can characterise it?

Definition 1. *A set S of Q-configurations is defined by ϕ if*

$$S = \left\{ M \in Q^{\mathbb{Z}^2} \middle| \mathfrak{M} \models \phi \right\}$$

Two formulas ϕ and ϕ' are equivalent iff they define the same set of configurations.

A set S is C-definable if it is defined by a formula $\phi \in C$.

It is easy to see that Example 1 is defined by the formula

$$\phi : \begin{cases} \forall x, \neg\left(P_{\color{red}\blacksquare}(x) \wedge P_{\color{cyan}\blacksquare}(\mathsf{East}(x))\right) \\ \forall x, \neg\left(P_{\color{red}\blacksquare}(x) \wedge P_{\color{cyan}\blacksquare}(\mathsf{North}(x))\right) \\ \forall x, \neg\left(P_{\color{cyan}\blacksquare}(x) \wedge P_{\color{red}\blacksquare}(\mathsf{East}(x)) \wedge P_{\color{cyan}\blacksquare}(\mathsf{North}(x))\right) \end{cases}$$

or equivalently by the formula

$$\phi' : \forall x, P_{\color{red}\blacksquare}(x) \iff \left(P_{\color{red}\blacksquare}(\mathsf{East}(x)) \wedge P_{\color{red}\blacksquare}(\mathsf{North}(x))\right)$$

We will see some variants of formula ϕ' appear in a few theorems below.

Example 2 is defined by the formula

$$\psi : \forall x, y, \left(P_{\blacksquare}(x) \wedge P_{\blacksquare}(y)\right) \implies x = y$$

Note that a definable set is always closed by shift (a shift between 2 configurations induces an isomorphism between corresponding structures). It is not always closed: The set of $\{{\color{lime}\blacksquare}, \square\}$-configurations defined by the formula ϕ : $\exists z, P_{\color{lime}\blacksquare}(z)$ contains all configurations except the all-white one, hence is not closed.

When we are dealing with MSO formulas, the following remark is useful: second-order quantifiers may be represented as projection operations on sets of configurations. We formalize now this notion.

If $\pi : Q_1 \mapsto Q_2$ is a projection and S is a set of Q_1-configurations, we define the two following operators:

$$
\begin{aligned}
E(\pi)(S) &= \left\{ M \in (Q_2)^{\mathbb{Z}^2} \middle| \exists N \in (Q_1)^{\mathbb{Z}^2}, \pi(N) = M \wedge N \in S \right\} \\
A(\pi)(S) &= \left\{ M \in (Q_2)^{\mathbb{Z}^2} \middle| \forall N \in (Q_1)^{\mathbb{Z}^2}, \pi(N) = M \implies N \in S \right\}
\end{aligned}
$$

Note that A is a dual of E, that is $A(\pi)(S) = {}^{c}E(\pi)({}^{c}S)$ where ^{c} represents complementation.

Proposition 1.

- *A set S of Q-configurations is EMSO-definable if and only if there exists a set S' of Q' configurations and a map $\pi : Q' \mapsto Q$ such that $S = E(\pi)(S')$ and S' is FO-definable.*
- *The class of MSO-definable sets is the closure of the class of FO-definable sets by the operators E and A.*

Proof (Sketch). We prove here only the first item.

- Let $\phi = \exists X, \psi$ be a EMSO formula that defines a set S of Q-configurations. Let $Q' = Q \times \{0, 1\}$ and π be the canonical projection from Q' to Q.

  Consider the formula ψ' obtained from ψ by replacing $X(t)$ by $\vee_{c\in Q} P_{(c,1)}(t)$ and $P_c(t)$ by $P_{(c,0)}(t) \vee P_{(c,1)}(t)$.

  Let S' be a set of Q' configurations defined by ψ'. Then is it clear that $S = E(\pi)(S')$. The generalization to more than one existential quantifier is straightforward.

- Let $S = E(\pi)(S')$ be a set of Q configurations, and S' FO-definable by the formula ϕ. Denote by $c_1 \ldots c_n$ the elements of Q'. Consider the formula ϕ' obtained from ϕ where each P_{c_i} is replaced by X_i. Let

$$
\psi = \exists X_1, \ldots, \exists X_n, \begin{cases} \forall z, \vee_i X_i(z) \\ \forall z, \wedge_{i\neq j}(\neg X_i(z) \vee \neg X_j(z)) \\ \forall z, \wedge_i (X_i z \implies P_{\pi(c_i)}(z)) \\ \phi' \end{cases}
$$

  Then ψ defines S. Note that the formula ψ constructed above is of the form $\exists X_1, \ldots, \exists X_n (\forall z, \psi'(z)) \wedge \phi'$. This will be important later. □

Second-order quantifications will then be regarded in this paper either as projections operators or sets quantifiers.

3. Hanf Locality Lemma and EMSO

The first-order logic has a property that makes it suitable to deal with tilings and configurations: it is local. This is illustrated by Hanf's lemma [20, 21, 22]. A square pattern of radius n is a pattern of domain $[-n, n] \times [n, n]$

Definition 2. *Two Q-configurations M and N are (n, k)-equivalent if for each Q-square pattern P of radius n:*

- *If P appears in M less than k times, then P appears the exact same number of times in M and in N*
- *If P appears in M more than k times, then P appears in N more than k times*

This notion is indeed an equivalence relation. Given n and k, it is clear that there is only finitely many equivalence classes for this relation.

The Hanf's local lemma can be formulated in our context as follows:

Theorem 2. *For every FO formula ϕ, there exists (n, k) such that*

if M and N are (n,k) equivalent, then $\mathfrak{M} \models \phi \iff \mathfrak{N} \models \phi$

Corollary 3. *Every FO-definable set is a (finite) union of some (n,k)-equivalence classes.*

This is theorem 3.3 in [14], stated for finite configurations. Lemma 3.5 in the same paper gives a proof of Hanf's Local Lemma in our context.

Given (P,k) we consider the set $S_{=k}(P)$ of all configurations such that the pattern P occurs exactly k times (k may be taken equal to 0). The set $S_{\geq k}(P)$ is the set of all configurations such that the pattern P occurs more than k times.

We may rephrase the preceding corollary as:

Corollary 4. *Every FO-definable set is a positive combination (i.e. unions and intersections) of some $S_{=k}(P)$ and some $S_{\geq k}(P)$*

Theorem 5. *Every EMSO-definable set can be defined by a formula ϕ of the form:*

$$\exists X_1,\dots,\exists X_n,\ (\forall z_1,\ \phi_1(z_1,X_1,\dots,X_n)) \\ \wedge (\exists z_1,\dots,\exists z_p,\ \phi_2(z_1\dots z_p,X_1,\dots,X_n)),$$

where ϕ_1 and ϕ_2 are quantifier-free formulas.

See [1, Corollary 4.1] or [23, Corollary 4.2] for a similar result. This result is an easy consequence of [24, Theorem 3.2] (see also the corrigendum). We include here a full proof.

Proof. Let C be the set of such formulas. We proceed in three steps:

- Every EMSO-definable set is the projection of a positive combination of some $S_{=k}(P)$ and $S_{\geq k}(P)$ (using prop. 1 and the preceding corollary)
- Every $S_{=}(P,k)$ (resp. $S_{\geq}(P,k)$) is C-definable
- C-definable sets are closed by (finite) union, intersection and projections.

C-definable sets are closed by projection using the equivalence of prop. 1 in the two directions, the note at the end of the proof and some easy formula equivalences. The same goes for intersection.

Now we prove that C-definable sets are closed by union. The difficulty is to ensure that we use only one universal quantifier. Let ϕ and ϕ' be two C-formulas defining sets S_1 and S_2. We can suppose that ϕ and ϕ' use the same numbers of second-order quantifiers and of first-order existential quantifiers.

Then the formula

$$\exists X, \exists X_1,\dots,\exists X_n, \forall z_1, \begin{cases} X(z_1) \iff X(\mathsf{North}(z_1)) \\ X(z_1) \iff X(\mathsf{East}(z_1)) \\ X(z_1) \implies \phi_1(z_1,X_1\dots X_n) \\ \neg X(z_1) \implies \phi_1'(z_1,X_1\dots X_n) \end{cases} \\ \wedge \exists z_1,\dots,\exists z_p \bigvee \begin{array}{l} X(z_1)\wedge\phi_2(z_1\dots z_p,X_1\dots X_n) \\ \neg X(z_1)\wedge\phi_2'(z_1\dots z_p,X_1\dots X_n) \end{array}$$

defines $S_1 \cup S_2$ (the disjunction is obtained through variable X which is forced to represent either the empty set or the whole plane $\mathbb{Z}^2$).

It is now sufficient to prove that a $S_{=k}(P)$ set (resp. a $S_{\geq k}(P)$ set) is definable by a C-formula. Let $\phi_P(z)$ be the quantifier-free formula such that $\phi_P(z)$ is true if and only if P appears at position z.

Then $S_{=k}(P)$ is definable by

$$\exists X_1\dots\exists X_k \exists A_1,\dots,\exists A_k, \forall x \begin{cases} \wedge_i A_i(x) \iff [A_i(\mathsf{North}(x))\wedge A_i(\mathsf{East}(x))] \\ \wedge_i X_i(x) \iff [A_i(x)\wedge\neg A_i(\mathsf{South}(x))\wedge\neg A_i(\mathsf{West}(x))] \\ \wedge_{i\neq j} X_i(x) \implies \neg X_j(x) \\ (\vee_i X_i(x)) \iff \phi_P(x) \end{cases} \\ \wedge \exists z_1,\dots,\exists z_k, X_1(z_1)\wedge\dots\wedge X_k(z_k)$$

The formula ensures indeed that A_i represents a quarter of the plane, X_i being a singleton representing the corner of that plane. If $k = 0$ this becomes $\forall x, \neg\phi_P(x)$. To obtain a formula for $S_{\geq k}(P)$, change the last $\iff$ to a $\implies$ in the formula. □

4. Characterization of Subshifts of Finite Type and Sofic Subshifts

4.1. Subshifts of Finite Type

We start by a characterization of subshifts of finite type (SFTs, i.e tilings). The problem with SFTs is that they are closed neither by projection nor by union. As a consequence, the corresponding class of formulas is not very interesting:

Theorem 6. *A set of configurations is a SFT if and only if it is defined by a formula of the form*

$$\forall z,\ \psi(z)$$

where ψ is quantifier-free.

Note that there is only one quantifier in this formula. Formulas with more than one universal quantifier do not always correspond to SFT: This is due to SFTs not being closed by union.

Proof. Let $P_1 \dots P_n$ be patterns. To each P_i we associate the quantifier-free formula $\phi_{P_i}(z)$ which is true if and only if P_i appears at the position z. Then the subshifts that forbids patterns $P_1 \dots P_n$ is defined by the formula:

$$\forall z, \neg\phi_{P_1(z)} \wedge \cdots \wedge \neg\phi_{P_n(z)}$$

Conversely, let ψ be a quantifier-free formula. Each term t_i in ψ is of the form $f_i(z)$ where f_i is some combination of the functions North, South, East and West, each f_i thus representing somehow some vector z_i ($f_i(z) = z + z_i$). Let Z be the collection of all vectors z_i that appear in the formula ψ. Now the fact that ψ is true at the position z only depends on the colors of the configurations in points $(z + z_1), \dots, (z + z_n)$, i.e. on the *pattern* of domain Z that occurs at position z. Let $\mathcal{P}$ be the set of patterns of domain Z that makes ψ false. Then the set S defined by ψ is the set of configurations where no patterns in $\mathcal{P}$ occurs, hence a SFT. □

4.2. Universal sentences

Due to the way subshifts are defined, universal quantifiers play an important role. We now ask the following question: what are the sets defined by universal formulas? First the following lemma shows that we can restrict to first-order when considering universal formulas.

Lemma 7. *Any universal formula is equivalent to a first-order universal formula.*

Proof. A universal formula is equivalent (through permutation of universal quantifiers) to a formula of the form

$$\forall x_1, \dots, x_p, \forall X_1, \dots, X_n,\ \Phi(X_1, \dots, X_n, x_1, \dots, x_p)$$

where Φ is quantifier-free. Consider the formula

$$\psi(X_1, \dots, X_{n-1}, x_1, \dots, x_p) \equiv \forall X_n, \Phi(X_1, \dots, X_n, x_1, \dots, x_p)$$

Let $\{t_1, \dots, t_k\}$ be the set of terms t such that $X_n(t)$ occurs in Φ. The idea is that the truth value of $\Phi(X_1, \dots, X_n, x_1, \dots, x_p)$ depends only on the value of X_n at positions represented by the (t_i). Depending on interpretations of the variables (x_i), interpretations of the terms (t_i) may be equal or not. We say an assignation $\rho : \{1, \dots, k\} \to \{0, 1\}$ is *sound* if $t_i = t_j \implies \rho(i) = \rho(j)$. Denote by $\phi_\rho(x_1, \dots, x_p)$ the quantifier-free formula expressing this condition:

$$\phi_\rho(x_1, \dots, x_p) \equiv \bigwedge_{\{(i,j):\rho(i)\neq\rho(j)\}} t_j \neq t_j.$$

Let ψ_ρ denote the formula $\Phi[X_n(t_i) \leftarrow \rho(i)]$ obtained from Φ be replacing each occurrence of $X_n(t_i)$ by the truth value $\rho(i)$ and this for each $i \in \{1, \dots, k\}$. For any fixed $x_1, \dots, x_p$, the truth value of $\forall X_n \Phi(X_1, \dots, X_n, x_1, \dots, x_p)$ is the same as the truth value of the conjonction of formulas ψ_ρ for all sound ρ. Hence, we get that $\psi(X_1, \dots, X_{n-1}, x_1, \dots, x_p)$ is equivalent to the following quantifier-free formula:

$$\bigwedge_{\rho:\{1,\dots,k\}\to\{0,1\}} \phi_\rho \implies \psi_\rho.$$

We can eliminate this way second order universal quantifiers one by one and the lemma follows. $\square$

For the rest of this section we focus on first-order universal formulas. The real difficulty is to treat the equality predicate (=). Without the equality (more precisely if all predicates and functions are only unary) any first-order universal formula is equivalent to a conjonction of formulas with only one quantifier and theorem 6 applies. The equality predicate intertwines the variables and makes thing a bit harder to prove. The reader might for example try to understand what the following formula exactly means:

$$\forall x, y, \left(P_{\square}(x) \wedge P_{\blacksquare}(\mathsf{East}(y))\right) \implies x = y$$

To understand it, we will prove an analog of Hanf's Lemma for universal sentences.

Definition 3. *Let (n, k) be integers, and M, N two Q-configurations. We say that $M \geq_{n,k} N$ if for each Q-square pattern P of radius less than n:*

- *If P appears in M exactly p times and $p \leq k$, then P appears less than p times in N*
- *(No condition is required if P appears in M more than k times)*

Note that M and N are (n, k) equivalent if and only if $M \geq_{n,k} N$ and $N \geq_{n,k} M$.

Theorem 8. *For every universal formula ϕ there exists (n, k) such that if $M \geq_{n,k} N$, then $\mathfrak{M} \models \phi \implies \mathfrak{N} \models \phi$*

Compare with definition 2 and theorem 2. Note that Gaifman's Theorem (a more refined version of Hanf's lemma) was generalized in [25] to existential sentences. We may use this result to obtain ours. This would however add some unnecessary complications.

Proof. We will translate the usual proof of Hanf's Local Lemma into our special case. We will try as much as possible to use the same notations as [21, sec. 2.4].

We first change the vocabulary and consider that $\mathsf{East}, \mathsf{West}, \mathsf{North}, \mathsf{South}$ are binary predicates rather than functions. Note that every universal formula will remain a universal formulas, albeit with more quantifiers.

Let introduce some notations. Let $S(r, a)$ be the set of all points at distance less than r of a. That is $S(r, a) = \{x : |x - a| \leq r\}$ where $|\cdot|$ is the Manhattan distance. Note that $S(r, a)$ contains $e_r = 2r^2 + 2r + 1$ points. Let $S(r, a_1 \dots a_p) = \cup_i S(r, a_i)$.

Let M and N be two Q-configurations. We say that $a_1 \dots a_p \in (\mathbb{Z}^2)^p$ and $b_1 \dots b_p \in (\mathbb{Z}^2)^p$ are k-isomorphic if there exists a bijective map f from $S(3^k, a_1 \dots a_p)$ to $S(3^k, b_1 \dots b_p)$ that preserves the relations, that is

- $x\ \mathsf{East}\ y \iff f(x)\ \mathsf{East}\ f(y)$
- $P_c(x) \iff P_c(f(x))$
- $f(a_i) = b_i$.

It is then clear that if $a_1 \dots a_p$ and $b_1 \dots b_p$ are 0-isomorphic, then we have $\mathfrak{M} \models \psi(a_1 \dots a_p) \iff \mathfrak{N} \models \psi(b_1 \dots b_p)$ whenever ψ is quantifier-free.

Now take a formula $\phi = \forall x_1 \dots x_n \psi(x_1 \dots x_n)$ where ψ is quantifier-free.

Let M and N such that $M \geq_{3^n, ne_{3^n}+1} N$.

We now prove by induction that

if $a_1 \dots a_p$ and $b_1 \dots b_p$ are $(n-p)$-isomorphic, then for all b_{p+1}, there exists a_{p+1} such that $a_1 \dots a_{p+1}$ and $b_1 \dots b_{p+1}$ are $(n-p-1)$-isomorphic.

- Case $p = 0$. Let $b_1 \in \mathbb{Z}^2$. Consider the pattern of radius 3^n centered around b_1 in N. This pattern appears in N, hence must appear in M at least one time. Take a_1 to be the center of this pattern.
- Case $p \mapsto p+1$. Let $a_1 \dots a_p$ and $b_1 \dots b_p$ be $n-p$ isomorphic. Let $b_{p+1} \in \mathbb{Z}^2$.
  - Case 1: $|b_{p+1} - b_i| \leq 2 \times 3^{n-p-1}$ for some b_i.
    In this case $S(3^{n-p-1}, b_{p+1}) \subseteq S(3^{n-p}, b_i)$. Hence by taking $a_{p+1} = f^{-1}(b_{p+1})$ where f is the bijective map involved in the $n-p$ isomorphism, it is clear that $a_1 \dots a_{p+1}$ and $b_1 \dots b_{p+1}$ are $n-p-1$ isomorphic.
  - Case 2: $\forall i, |b_{p+1} - b_i| > 2 \times 3^{n-p-1}$. In this case for every i, $S(3^{n-p-1}, b_{p+1}) \cap B(3^{n-p-1}, b_i) = \emptyset$.
    Consider the pattern P of radius 3^{n-p-1} centered around b_{p+1}.
    This pattern appears α times inside $S(2 \times 3^{n-p-1}, b_1 \dots b_p)$ where $\alpha \leq pe_{2\times 3^{n-p-1}}$. P appears at least $\alpha + 1$ times in N and $\alpha + 1 \leq ne_{3^n} + 1$ hence must appears at least $\alpha + 1$ times in M. As it appears the same amount of time in $S(2\times 3^{n-p-1}, b_1 \dots b_p)$ and $S(2\times 3^{n-p-1}, a_1 \dots a_p)$ (by $n-p$ isomorphism), it must appear somewhere else, say centered in a_{p+1}. This a_{p+1} is not inside $S(3^{n-p-1}, a_1 \dots a_p)$ because otherwise it would be the center of an occurrence of pattern P inside $S(2 \times 3^{n-p-1}, a_1 \dots a_p)$. As a consequence, $a_1 \dots a_{p+1}$ and $b_1 \dots b_{p+1}$ are $n-p-1$ isomorphic.

Now suppose that $\mathfrak{M} \models \phi$. Take $b_1 \dots b_n \in \mathbb{Z}^2$. There exists $a_1 \dots a_n$ such that $a_1 \dots a_n$ and $b_1 \dots b_n$ are 0-isomorphic. As $\mathfrak{M} \models \phi$ the quantifier-free formula $\psi(a_1 \dots a_n)$ is true in $\mathfrak{M}$. As a consequence $\psi(b_1 \dots b_n)$ is true in $\mathfrak{N}$. As this is true for all $b_1 \dots b_n$ we obtain $\mathfrak{N} \models \phi$. □

Given (P, k) we consider the set $S_{\leq k}(P)$ of all configurations such that the pattern P occurs at most k times (k may be taken equal to 0)

Corollary 9. *A set is definable by a universal formula if and only if it is a positive combination (i.e. unions and intersections) of some $S_{\leq k}(P)$.*

Compare to corollary 4.

Proof. Let C be the class of all universal formulas. It is clear that the set of C-defined formulas is closed under intersection and unions.

Now $S_{\leq k}(P)$ is defined by

$$\forall x_1 \dots x_{k+1}, \phi_P(x_1) \wedge \dots \wedge \phi_P(x_{k+1}) \implies \bigvee_{i \neq j} x_i = x_j$$

For $k = 0$, this becomes $\forall x, \neg\phi_P(x)$. Hence, every positive combination of some $S_{\leq k}(P)$ is C-definable.

Conversely, let ϕ be a universal formula and S the set it defines. Let (n, k) be as in the theorem.

For each configuration $M \in S$ and P a pattern of radius less than n, denote $\phi_M(P)$ the number of times P appears in M with the convention than $\phi_M(P) = \infty$ if P appears more than k times in M.

Consider the set

$$S_M = \bigcap_{\substack{P|\phi_M(P)\neq\infty,\\ \text{radius}(P)\leq n}} S_{\leq \phi_M(P)}(P)$$

From the hypothesis on (n, k), we have $S_M \subseteq S$. It is then easy to see that $S = \cup_M S_M$ where the union is actually finite (two configurations that are (n, k)-equivalent give the same S_M). □

4.3. Sofic subshifts

Using the previous corollary, we are now able to give a characterisation of sofic subshifts:

Theorem 10. *A set S is a sofic subshift if and only if it is definable by a formula of the form*

$$\exists X_1, \ldots, \exists X_n, \forall z_1, \ldots, \forall z_p,\ \psi(X_1, \ldots, X_n, z_1 \ldots z_p)$$

where ψ is quantifier-free. Moreover, any such formula is equivalent to a formula of the same form but with a single universal quantifier ($p = 1$).

See [19] for a different proof that eliminates equality predicates one by one.

Proof. Let C be the clas of all formulas of the form

$$\exists X_1, \ldots, \exists X_n, \forall z \psi(X_1, \ldots, X_n, z)$$

where ψ is quantifier-free. With the help of theorem 6 and proposition 1, is is quite clear that C-defined sets are exactly sofic subshifts.

Let $\mathcal{D}$ be the class of all formulas of the form

$$\exists X_1, \ldots, \exists X_n, \forall z_1 \ldots z_p \psi(X_1, \ldots, X_n, z_1 \ldots z_p)$$

where ψ is quantifier-free. The previous remark states that sofic subshifts are $\mathcal{D}$-defined.

Now we prove that $\mathcal{D}$-defined sets are sofic subshifts. Using (the proof of) proposition 1, and the fact that sofic subshifts are closed under projection, it is sufficient to prove that universal formulas define sofic subshifts. Using corollary 9 and the fact that sofic subshifts are closed under union and projections, it is sufficient to prove that every $S_{\leq k}(P)$ is sofic.

Now $S_{\leq k}(P)$ is defined by

$$\phi : \exists S_1 \ldots S_k \begin{cases} & \Psi_i \\ \forall x, \vee_i S_i(x) & \iff \phi_P(x) \end{cases}$$

where Ψ_i expresses that S_i has at most one element and is defined as follows:

$$\Psi_i \stackrel{def}{=} \exists A, \forall x \begin{cases} & A(x) \iff A(\mathrm{North}(x)) \wedge A(\mathrm{East}(x)) \\ S_i(x) \iff & A(x) \wedge \neg A(\mathrm{South}(x)) \wedge \neg A(\mathrm{West}(x)) \end{cases}$$

Now with some light rewriting we can transform ϕ into a formula of the class C, which proves that $S_{\leq k}(P)$ is C-definable, hence sofic. □

5. (E)MSO-definable subshifts

5.1. Separation result

Theorems 5 and 10 above suggest that EMSO-definable subshifts are not necessarily sofic. We will show in this section that the set of EMSO-definable subshifts is indeed strictly larger than the set of sofic subshifts. The proof is based on the analysis of the computational complexity of forbidden languages. It is well-known that sofic subshifts have a recursively enumerable forbidden language. The following theorem shows that the forbidden language of an MSO-definable subshift can be arbitrarily high in the arithmetical hierarchy.

This is not surprising since arbitrary Turing computation can be defined via first order formulas (using tilesets) and second order quantifiers can be used to simulate quantification of the arithmetical hierarchy. However, some care must be taken to ensure that the set of configurations obtained is a subshift.

Theorem 11. *Let E be an arithmetical set. Then there is an MSO-definable subshift with forbidden language $\mathcal{F}$ such that E reduces to $\mathcal{F}$ (for many-one reduction).*

Proof (sketch). Suppose that the complement of E is defined as the set of integers m such that:

$$\exists x_1, \forall x_2, \ldots, \exists/\forall x_n, R(m, x_1, \ldots, x_n)$$

where R is a recursive relation. We first build a formula ϕ defining the set of configurations representing a successful computation of R on some input $m, x_1, \ldots, x_n$. Consider 3 colors c_l, c and c_r and additional second order variables $X_1, \ldots, X_n$ and $S_1, \ldots, S_n$. The input $(m, x_1, \ldots, x_n)$ to the computation is encoded in unary on an horizontal segment using colors c_l and c_r and variables S_i as separators, precisely: first an occurrence of c_l then m occurrences of c, then an occurrence of c_r and, for each successive $1 \leq i \leq n$, x_i positions in X_i before a position of S_i. Let ϕ_1 be the FO formula expressing the following:

1. there is exactly 1 occurrence of c_l and the same for c_r and all S_i are singletons;
2. starting from an occurrence c_l and going east until reaching S_n, the only possible successions of states are those forming a valid input as explained above.

Now, the computation of R on any input encoded as above can be simulated via tiling constraints in the usual way. Consider sufficiently many new second order variables $Y_1, \ldots, Y_p$ to handle the computation and let ϕ_2 be the FO formula expressing that:

1. a valid computation starts at the north of an occurrence of c_l;
2. there is exactly one occurrence of the halting state (represented by some Y_i) in the whole configuration.

We define ϕ by:

$$\exists X_1, \forall X_2, \ldots, \exists/\forall X_n, \exists S_1, \ldots, \exists S_n, \exists Y_1, \ldots, \exists Y_p, \phi_1 \wedge \phi_2.$$

Finally let ψ be the following FO formula: $(\forall z, \neg P_{c_l}) \vee (\forall z, \neg P_{c_r})$. Let X be the set defined by $\phi \vee \psi$. By construction, a finite (unidimensional) pattern of the form $c_l c^m c_r$ appears in some configuration of X if and only if $m \notin E$. Therefore E is many-one reducible to the forbidden language of X.

To conclude the proof it is sufficient to check that X is closed. To see this, consider a sequence $(C_n)_n$ of configurations of X converging to some configuration C. C has at most one occurrence of c_l and one occurrence of c_r. If one of these two states does not occur in C then $C \in X$ since ψ is verified. If, conversely, both c_l and c_r occur (once each) then any pattern containing both occurrences also occurs in some configuration C_n verifying ϕ. But ϕ is such that any modification outside the segment between c_l and c_r in C_n does not change the fact that ϕ is satisfied provided no new c_l and c_r colors are added. Therefore ϕ is also satisfied by C and $C \in X$. □

The theorem gives the claimed separation result for subshifts of EMSO.

Corollary 12. *There are EMSO-definable subshifts which are not sofic.*

Proof. In the previous theorem, choose E, to be the complement of the set of integers m for which there is x such that machine m halts on empty input in less than x steps. E is not recursively enumerable and, using the construction of the proof above, it is reducible to the forbidden language of an EMSO-definable subshift. □

5.2. Definability of MSO-subshifts

As we saw before, sets defined by MSO-formulas are not always subshifts. We will try in this section to find a fragment of MSO that contains only subshifts and contain all of them. This fragment is somewhat ad hoc. Finding a more reasonable fragment is an interesting open question.

We first begin by a definition

Definition 4.

$$fin(S) : \exists A, \exists B \begin{cases} \forall x, A(x) \iff A(\mathit{North}(x)) \wedge A(\mathit{East}(x)) \\ \forall x, B(x) \iff A(\mathit{South}(x)) \wedge A(\mathit{West}(x)) \\ \exists x, A(x) \wedge \neg A(\mathit{South}(x)) \wedge \neg A(\mathit{West}(x)) \\ \exists x, B(x) \wedge \neg B(\mathit{North}(x)) \wedge \neg B(\mathit{East}(x)) \\ \forall x, S(x) \implies A(x) \wedge B(x) \end{cases}$$

It is easy to prove that $fin(S)$ is true if and only if S is finite (there are finitely many x such that $S(x)$). Indeed A and B represent quarter of planes, and S must be contained in the square delimited by the two quarter of planes. Any other formula true only if S is finite would work in the following

Theorem 13. *Let S be a MSO-definable set. Then S is a subshift if and only if it is definable by a formula of the form*

$$\forall S, fin(S) \implies \exists B_1 \dots B_k, \psi(S, B_1 \dots B_k) \wedge \forall x_1 \dots x_n S(x_1) \wedge \dots S(x_p) \implies \theta(S, B_1 \dots B_k, x_1 \dots x_p)$$

where

- *ψ is any MSO-formula not containing the predicates P_c.*
- *θ is quantifier-free.*

Note that this formula can be written more concisely as

$$\forall^{fin} S, \exists \overline{B} \psi(S, \overline{B}) \wedge \forall \overline{x} \in S^p, \theta(S, \overline{B}, \overline{x})$$

Proof. First we prove that such a formula ϕ defines a subshift X. For this, we prove that the set X is closed. Consider a sequence $M_1 \dots M_n \dots$ of configurations of X converging to some configuration M. We must prove that $M \in X$.

Let S be a finite set. Now consider the formula θ. As it is quantifier-free, it is local: the value of $\theta(S, B_1 \dots B_k, x_1 \dots x_n)$ depends only of what happens around $x_1 \dots x_n$. As each $x_1 \dots x_n$ must be in S, there exists a finite $S' \supset S$ such that the value of $\forall x_1 \in S \dots x_n \in S, \theta(S, B_1 \dots B_k, x_1 \dots x_n)$ depends only of the value of the predicates S, P_c and B_i on S'.

Now M_i converges to M. This means that there exists p such that M_p and M coincides on S'. For this M_p, there exists some $B_1 \dots B_k$ such that we have $\mathfrak{M}_p \models \psi(S, B_1 \dots B_k) \wedge \forall x_1 \in S \dots \forall x_p \in S, \theta(S, B_1 \dots B_k, x_1 \dots x_n)$. Then this formula is also true on $\mathfrak{M}$ (Note indeed that $\psi(S, B_1 \dots B_k)$ does not depend on the configuration).

Hence we have found for every S some B_i that makes the formula true, that is we have proven $\mathfrak{M} \models \phi$. Therefore X is closed, hence a subshift.

Now let X be a MSO-definable subshift. X is defined by a formula ϕ. Change each P_c in ϕ by a predicate B_c to obtain ψ_1. Define

$$\psi(\overline{B}) = \forall x \left(\bigvee_c B_c(x) \right) \wedge \left(\bigwedge_{c \neq c'} \neg (B_c(x) \wedge B_{c'}(x)) \right) \wedge \psi_1(\overline{B})$$

Then X is defined by

$$\phi : \forall^{fin} S, \exists \overline{B} \psi(\overline{B}) \wedge \forall x \in S, \bigwedge_c (B_c(x) \iff P_c(x))$$

Indeed M satisfies ϕ and only if every pattern of M is a pattern in some configuration of X. □

6. A Characterization of EMSO

EMSO-definable sets are projections of FO-definable sets (proposition 1). Besides, sofic subshifts are projections of subshifts of finite type (or tilings). Previous results show that the correspondence sofic↔EMSO fails. However, we will show in this section how EMSO can be characterized through projections of "locally checkable" configurations.

Corollary 4 expresses that FO-definable sets are essentially captured by counting occurrences of patterns up to some value. The key idea in the following is that this counting can be achieved by local checkings (equivalently, by tiling constraints), provided it is limited to a finite and explicitly delimited region. This idea was successfully used in [14] in the context of picture languages: pictures are rectangular finite patterns with a border made explicit using a special state (which occurs all along the border and nowhere else). We will proceed here quite differently. Instead of putting special states on borders of some rectangular zone, we will simply require that two special subsets of states Q_0 and Q_1 are present in the configuration: we call a (Q_0, Q_1)-*marked configuration* any configuration that contains both a color $q \in Q_0$ and some color $q' \in Q_1$ somewhere. By extension, given a subshift Σ over Q and two subsets $Q_0 \subseteq Q$ and $Q_1 \subseteq Q$, the *doubly-marked set* Σ_{Q_0,Q_1} is the set of (Q_0, Q_1)-marked configurations of Σ. Finally, a *doubly-marked set of finite type* is a set Σ_{Q_0,Q_1} for some SFT Σ and some Q_0, Q_1.

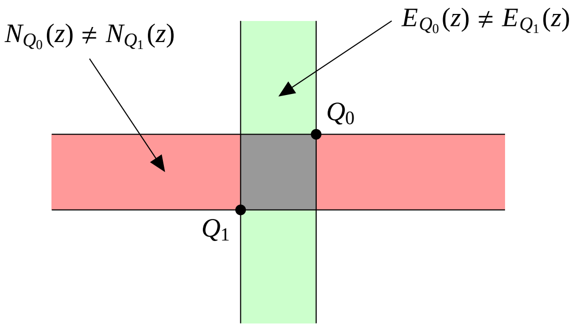

Figure 4: The rectangular zone in dark gray defined by predicate $Z(z)$.

Lemma 14. *For any finite pattern P and any $k \geq 0$, $S_{=k}(P)$ is the projection of some doubly-marked set of finite type. The same result holds for $S_{\geq k}(P)$.*

Moreover, any positive combination (union and intersection) of projections of doubly-marked sets of finite type is also the projection of some doubly-marked sets of finite type.

Proof (sketch). We consider some base alphabet Q, some pattern P and some $k \geq 0$. We will build a doubly-marked set of finite type over alphabet $Q' = Q \times Q_+$ and then project back on Q. Q_+ is itself a product of different layers. The first layer can take values $\{0, 1, 2\}$ and is devoted to the definition of the marker subsets Q_0 and Q_1: a state is in Q_i for $i \in \{0, 1\}$ if and only if its value on the layer is i.

We first show how to convert the apparition in a configuration of two marked positions, by Q_0 and Q_1, into a locally identifiable rectangular zone. The zone is defined by two opposite corners corresponding to an occurrence of some state of Q_0 and Q_1 respectively. This can be done using only finite type constraints as follows. By adding a new layer of states, one can ensure that there is a unique occurrence of a state of Q_0 and maintain everywhere the following information:

1. $N_{Q_0}(z) \equiv$ the position z is at the north of the (unique) occurrence of a state from Q_0,
2. $E_{Q_0}(z) \equiv$ the position z is at the east of the occurrence of a state from Q_0.

The same can be done for Q_1. From that, the membership to the rectangular zone is defined at any position z by the following predicate (see figure 11):

$$Z(z) \equiv N_{Q_0}(z) \neq N_{Q_1}(z) \wedge E_{Q_0}(z) \neq E_{Q_1}(z).$$

We can also define locally the border of the zone: precisely, cells not in the zone but adjacent to it. Now define $P(z)$ to be true if and only if z is the lower-left position in an occurrence of the pattern P. We add k new layers, each one storing (among other things) a predicate $C_i(z)$ verifying

$$C_i(z) \Rightarrow Z(z) \wedge P(z) \wedge \bigwedge_{j \neq i} \neg C_j(z).$$

Moreover, on each layer i, we enforce that exactly 1 position z verifies $C_i(z)$: this can be done by maintaining north/south and east/west tags (as for Q_0 above) and requiring that the north (resp. south) border of the rectangular zone sees only the north (resp. south) tag and the same for east/west. Finally, we add the constraint:

$$P(z) \wedge Z(z) \Rightarrow \bigvee_i C_i$$

expressing that each occurrence of P in the zone mut be "marked" by some C_i. Hence, the only admissible (Q_0, Q_1)-marked configurations are those whose rectangular zone contains exactly k occurrences of pattern P. We thus obtain exactly $S_{\geq k}(P)$ after projection. To obtain $S_{=k}(P)$, it suffices to add the constraint:

$$P(z) \Rightarrow Z(z)$$

in order to forbid occurrences of P outside the rectangular zone.

To conclude the proof we show that finite unions or intersections of projections of doubly-marked sets of finite type are also projections of doubly-marked sets of finite type. Consider two SFT X over Q and Y over Q' and two pairs of marker subsets $Q_0, Q_1 \subseteq Q$ and $Q'_0, Q'_1 \subseteq Q'$. Let $\pi_1 : Q \to A$ and $\pi_2 : Q' \to A$ be two projections.

First, for the case of union, we can suppose (up to renaming of states) that Q and Q' are disjoint and define the SFT Σ over alphabet $Q \cup Q'$ as follows:

- 2 adjacent positions must be both in Q or both in Q';
- any pattern forbidden in X or Y is forbidden in Σ.

Clearly, $\pi(\Sigma_{Q_0 \cup Q'_0, Q_1 \cup Q'_1}) = \pi_1(X_{Q_0,Q_1}) \cup \pi_2(Y_{Q'_0,Q'_1})$ where $\pi(q)$ is $\pi_1(q)$ when $q \in Q$ and $\pi_2(q)$ else.

Now, for intersections, consider the SFT Σ over the fiber product

$$Q_\times = \{(q, q') \in Q \times Q' | \pi_1(q) = \pi_2(q')\}$$

and defined as follows: a pattern is forbidden if its projection on the component Q (resp. Q') is forbidden in X (resp. Y);

If we define π as π_1 applied to the Q-component of states, and if E is the set of configuration of Σ such that states from Q_0 and Q_1 appear on the first component and states from Q'_0 and Q'_1 appear on the second one, then we have:

$$\pi(E) = \pi_1(X_{Q_0,Q_1}) \cup \pi_2(Y_{Q'_0,Q'_1}).$$

To conclude the proof, it is sufficient to obtain E as the projection of some doubly-marked set of finite type. This can be done starting from Σ and adding a new component of states whose behaviour is to define a zone from two markers (as in the first part of this proof) and check that the zone contains occurrences of Q_0, Q_1, Q'_0 and Q'_1 in the appropriate components. □

Theorem 15. *A set is EMSO-definable if and only if it is the projection of a doubly-marked set of finite type.*

Proof. First, a doubly-marked set of finite type is an FO-definable set because SFT are FO-definable (theorem 6) and the restriction to doubly-marked configurations can be expressed through a simple existential FO formula. Thus the projection of a doubly-marked set of finite type is EMSO-definable.

The opposite direction follows immediately from proposition 1 and corollary 4 and the lemma above. □

At this point, one could wonder whether considering simply-marked set of finite type is sufficient to capture EMSO via projections. In fact the presence of 2 markers is necessary in the above theorem: considering the set Σ_{Q_0,Q_1} where Σ is the full shift $Q^{\mathbb{Z}^2}$ and Q_0 and Q_1 are distinct singleton subsets of Q, a simple compactness argument allows to show that it is not the projection of any simply-marked set of finite type.

7. Open Problems

- Is the second order alternation hierarchy strict for MSO (considering our model-theoretic equivalence)?
- One can prove that theorem 6 also holds for formulas of the form:

  $$\forall X_1 \ldots \forall X_n, \forall z, \psi(z, X_1 \ldots X_n)$$

  where ψ is quantifier-free. Hence, adding universal second-order quantifiers does not increase the expression power of formulas of theorem 6. More generally, let C be the class of formulas of the form

  $$\forall X_1, \exists X_2, \ldots, \forall/\exists X_n, \forall z_1, \ldots, \forall z_p, \phi(X_1, \ldots, X_n, z_1, \ldots, z_p).$$

  One can check that any formula in C defines a subshift. Is the second-order quantifiers alternation hierarchy strict in C? On the contrary, do all formulas in C represent sofic subshifts ?

References

[1] W. Thomas, Handbook of Formal Languages, Vol. 3. Beyond Words, Springer, 1997, Ch. Languages, Automata, and Logic.

[2] O. Matz, N. Schweikardt, Expressive power of monadic logics on words, trees, pictures, and graphs, in: E. G. J. Flum, T. Wilke (Eds.), Logic and Automata: History and Perspectives, Texts in Logic and Games, Amsterdam University Press, 2007, pp. 531–552.

[3] H. Wang, Proving theorems by pattern recognition ii, Bell system technical journal 40 (1961) 1–41.

[4] R. Berger, The undecidability of the domino problem, Memoirs American Mathematical Society 66 (1966) 1966.

[5] Y. G. Ergon Börger, Erich Grädel, The classical decision problem, Springer-Verlag Telos, 1996.

[6] D. Kuske, M. Lohrey, Logical aspects of Cayley-graphs: the group case, Annals of Pure and Applied Logic 131 (1-3) (2005) 263–286.

[7] D. Seese, The structure of the models of decidable monadic theories of graphs, Annals of pure and applied logic 53 (2) (1991) 169–195.

[8] J. Makowsky, On some conjectures connected with complete sentences, Fund. Math. 81 (1974) 193–202.

[9] B. Poizat, Une théorie finiement axiomatisable et superstable, Groupe d'étude des théories stables 3 (1) (1980-1982) 1–9.

[10] R. M. Robinson, Undecidability and Nonperiodicity for Tilings of the Plane, Inventiones Math. 12.

[11] F. Oger, Algebraic and model-theoretic properties of tilings, Theoretical computer science 319 (2004) 103–126.

[12] A. Ballier, E. Jeandel, Tilings and model theory, First Symposium on Cellular Automata Journées Automates Cellulaires.

[13] D. Giammarresi, A. Restivo, Handbook of Formal Languages, Vol. 3. Beyond Words, Springer, 1997, Ch. Two-Dimensional Languages.

[14] D. Giammarresi, A. Restivo, S. Seibert, W. Thomas, Monadic Second-Order Logic over Rectangular Pictures and Recognizability by Tiling Systems, Information and Computation 125 (1996) 32–45.

[15] B. Weiss, Subshifts of finite type and sofic systems, Monatshefte für Mathematik 77 (1973) 462–474.

[16] M. Anselmo, N. Jonoska, M. Madonia, Framed Versus Unframed Two-Dimensional Languages, in: 35th Conference on Current Trends in Theory and Practice of Computer Science, SOFSEM, 2009, pp. 79–92.

[17] J.-H. Altenbernd, W. Thomas, S. Wohrle, Tiling systems over infinite pictures and their acceptance conditions, in: Developments in Language Theory, 2003.

[18] D. Harel, Recurring Dominoes: Making the Highly Undecidable Highly Understandable, Annals of Discrete Mathematics 24 (1985) 51–72.

[19] E. Jeandel, G. Theyssier, Subshifts, Languages and Logic, in: Developments in Language Theory (DLT), no. 5583 in LNCS, Springer, 2009, pp. 288–299.

[20] W. Hanf, Model-theoretic methods in the study of elementary logic, in: Symposium in the Theory of Models, 1965, pp. 132–145.

[21] H.-D. Ebbinghaus, J. Flum, Finite Model Theory, Springer Monographs in Mathematics, Springer, Berlin, 1995.

[22] L. Libkin, Elements of Finite Model Theory, Texts in Theoretical Computer Science, an EATCS Series, Springer Verlag, 2004.

[23] W. Thomas, On logics, tilings, and automata, in: S. Berlin (Ed.), Proceedings of the 18th ICALP, Vol. 510 of LNCS, 1991, pp. 441–454.

[24] T. Schwentick, K. Barthelmann, Local Normal Forms for First-Order Logic with Applications to Games and Automata, Discrete Mathematics and Theoretical Computer Science 3 (1999) 109–124.

[25] M. Grohe, S. Wöhrle, An existential locality theorem, Annals of Pure and Applied Logic 129 (1-3) (2004) 131–148.